Submitted to: Neural Computation

Correlated Firing Improves Stimulus Discrimination

Garrett T. Kenyon[1,4], James Theiler[2,4], John S. George[1,4], Bryan J. Travis[3,4], David W. Marshak[5]

[1]Physics Division (P-21), [2]Non-Proliferation and International Security (NIS-2), [3]Environmental and Earth Sciences (EES-6), [4]Los Alamos National Laboratory, Los Alamos, NM, 87545; [5]Department of Neurobiology and Anatomy, University of Texas Medical School, Houston, TX, 77030

Abbreviated Title: Correlated Inputs to Model Neurons

Corresponding Author:
Garrett T. Kenyon
P-21 Biological and Quantum Physics, MS D454
Los Alamos National Laboratory
Los Alamos, NM 87544
phone: 505-667-1900
fax: 505-665-4507
email: gkenyon@lanl.gov

Acknowledgements: The authors wish to acknowledge useful discussions with Rob Smith. This work was supported by the Department of Energy Office of Nonproliferation Research and Engineering, the MIND Institute for Functional Brain Imaging, National Eye Institute #EY06472, the National Institute of Neurological Disease and Stroke #NS38310, and by the Lab Directed Research and Development Program at the Los Alamos National Laboratory.

Keywords: synchrony, gamma oscillations, retinal ganglion cell, ideal observer, rate code, computer model

Abstract

Synchronous oscillations impose strict lower limits on the amount of noise-reduction achievable by averaging individual neuron firing rates. Here, we show that the loss of rate-coded information due to spatiotemporal correlations, such as synchrony, can be overcome by encoding the same information in the correlations themselves. A retinal model was used to generate artificial spike trains whose spatiotemporal correlations were modulated by stimulus intensity in a manner consistent with experimental data. Pooled multi-unit spike trains were summed into a threshold detector whose output was classified by an ideal observer. For a threshold detector with short summation times, realistic spatiotemporal correlations yielded superior discrimination of stimulus intensity compared to independent Poisson controls. Even for summation times too long to resolve synchronous inputs, temporal correlations still contributed to improved discrimination by reducing spike count variability. Our results show that by directly encoding stimulus properties, spatiotemporal correlations in pooled multi-unit data mediate improved or equivalent information transmission compared to statistically independent rate-matched controls.

Introduction

Neurons represent sensory information as changes in their individual firing rates. In order to obtain reliable estimates of rate-coded information on physiological time scales, however, from 10's to 100's of msec, it may be necessary to average over ensembles of similarly activated cells. While this requirement is consistent with the high levels of convergence onto most central neurons, averaging only yields a more accurate estimate of rate-coded information to the extent that the input spike trains are statistically independent [25,33,34]. When the same input spike trains are instead strongly correlated, it is no longer possible to extract the same amount of rate-coded information over equivalent time intervals, as fluctuations in the number of spikes arising from individual cells will no longer average out across the population of input fibers. Nonetheless, correlated firing is ubiquitous in the nervous system, often associated with coherent oscillations in the gamma frequency band (>50 Hz) that synchronizes activity both within and between brain areas [12,35]. Here, we use 'spatial' correlations to refer to synchrony between cells at separate locations and 'temporal' correlations to refer to periodic or other reproducible firing patterns that can be measured at a single location.

It is important to understand what consequences strong spatiotemporal correlations might have for processing rate-coded information. One possibility is that strong spatiotemporal correlations simply impose an upper limit on the amount of rate-coded information a population of neurons can represent in a given unit of time [25,33,34]. From this perspective, firing correlations are an inevitable but undesirable consequence of the massive interconnectivity of neural circuits but otherwise serve no significant information processing function. A second possibility is that correlations instead play an essential information processing role, a view supported by studies showing that coherent oscillations may be involved in a variety of cognitive processes, including attention [11], perception [37], top-down priming [9], and feature integration

[12,35]. Here, we present evidence that these two viewpoints can be reconciled, namely, that the information encoded by the correlations themselves can overcome the attendant loss of rate-coded information.

Using a realistic model of how visual information is encoded in the spike trains of retinal output neurons, called ganglion cells, we show that the same stimulus properties can be encoded both by the firing rates of individual cells and by their mutual correlations. We further show that such hybrid rate/correlation codes can accomplish equal or superior reliability of information transmission when extracted by a threshold detection process. Information transmission was assessed by quantifying the performance of an ideal observer required to discriminate between different stimulus intensities based on the output of a threshold detector. When driven by correlated inputs in which the stimulus intensity was encoded by both the mean firing rate and by the level of synchrony, threshold elements mediated equal or superior stimulus discrimination than when driven by statistically independent Poisson generators that produced, on average, the same number of spikes. Threshold detectors with short integration times and low background event rates (i.e. coincidence detectors) always extracted more information from spike trains with realistic correlations than from Poisson controls, as the stimulus information encoded by the level of synchrony more than compensated for the attendant loss of rate-coded information. Neurons in the LGN and visual cortex are preferentially sensitive to synchronous inputs [1,41] and may serve to detect coincidences in retinal input. Surprisingly, even threshold detectors with long integration times, which were thus unable to resolve synchronous inputs, still mediated equal or superior stimulus discrimination when driven by correlated inputs compared to Poisson controls. Because temporal correlations became stronger as a function of the stimulus intensity, the loss of rate-coded information was countered by an increased reliability in the total number of spikes.

Our results are not in conflict with information theory, which stipulates that the capacity of a channel is always maximized when the inputs are independent. In the present analysis, inputs are always pooled into a single, multi-unit measure. The only benefit of statistical independence is therefore an improvement in signal to noise. Our results, however, suggest that the reduction in noise can be outweighed by the information carried by the correlations themselves. We propose that the nervous system may employ a hybrid encoding strategy that utilizes the enhanced representational capacity of spatiotemporal correlations without producing substantial losses in rate-coded information.

Methods

Retinal Model

Artificial spike trains with realistic spatiotemporal correlations were generated by a retinal model (fig 1), organized as a 32x32 array with wrap-around boundary conditions containing 5 distinct cell types: Bipolar cells (BP), small amacrine cells (SA), large amacrine cells (LA), poly-axonal amacrine cells (PA), and alpha ganglion cells (GC). All cell types were modeled as single compartment, RC circuit elements obeying a first order differential equation that can be written efficiently in terms of matrix multiplications:

$$\dot{\vec{V}}^{(k)} = -\frac{1}{\tau^{(k)}}\left(\vec{V}^{(k)} - \sum_{k'} \vec{W}^{(k,k')} \cdot f^{(k,k')}\left(\vec{V}^{(k')}\right) \cdot \vec{W}^{(k,k')^T} - b^{(k)} - \vec{L}^{(k)}\right), \qquad (1)$$

where $\vec{V}^{(k)}$ is a 2-D array denoting the normalized membrane potentials of all cells of type k, $(1 \le k \le 5)$, $\tau^{(k)}$ is the time constant, $b^{(k)}$ is a bias current for setting the resting potential, $\vec{L}^{(k)}$ is an external input intended to represent light stimulation, $\vec{W}^{(k,k')}$ gives the connection strengths between presynaptic $\{k'\}$ and postsynaptic $\{k\}$ cell types as a function of their separation along one direction, defined here as 'vertical', $\vec{W}^{(k,k')^T}$ gives the same information as a function of separation along the perpendicular direction, defined here as 'horizontal', and the functions $f^{(k,k')}$ give the associated input-output relations for the indicated pre- and post-synaptic cell types, detailed below. The output of the axon-mediated inhibition was delayed by 2 msec, except for the axonal connections onto the axon-bearing amacrine cells, which was delayed for 1 msec. All other synaptic interactions were delayed by one time step, which equaled 1 msec. All equations were integrated in Matlab using a direct Euler method.

The input-output function for gap junctions was given by the identity:

$$f^{(k,k')}\left(\vec{V}^{(k')}\right) = \vec{V}^{(k')}, \qquad (2)$$

where the dependence on the presynaptic potential has been absorbed into the definition of $\tau^{(k)}$. This is possible because both the decay term in equation 1 and the omitted dependence on the presynaptic potential in equation 2 depend linearly on $\vec{V}^{(k)}$, allowing the coefficients to be combined. The input-output function for non-spiking synapses was constructed by comparing, on each time step, a random number with a Fermi-function:

$$f^{(k,k')}\left(\vec{V}^{(k')}\right)=\theta\left(\left[\frac{1}{1+\exp\left(-\alpha\vec{V}^{(k')}\right)}\right]-r\right), \tag{3}$$

where α sets the gain (equal to 4 for all non-spiking synapses), r is a uniform random deviate equally likely to take any real value between 0 and 1, and θ is a step function, $\theta(x)=1,\ x\geq 0$; $\theta(x)=0,\ x<0$.

Lastly, the input-output relation used for spiking synapses was:

$$f^{(k,k')}\left(\vec{V}^{(k')}\right)=\theta\left(\vec{V}^{(k')}\right). \tag{4}$$

A modified integrate-and-fire mechanism was used to model spike generation. A positive pulse (amplitude = 10.0) was delivered to the cell on the time step after the membrane potential crossed threshold, followed by a negative pulse (amplitude = -10.0) on the subsequent time step. This resulted in a 1 msec action potential that also produced impulse responses in electrically coupled cells, an important element of the circuit dynamics. The bias current, *b*, was incremented by -0.5 following each spike, and then decayed back to the resting value with the time constant of the cell, adding to the relative refractory period. There was in addition an absolute refractory period of 1 msec.

Along both the horizontal and vertical directions, synaptic strengths fell off as Gaussian functions of the distance between the pre- and post-synaptic cells. For a given horizontal separation, the horizontal weight factor was determined by a Gaussian function of the following form:

$$W^{(k,k')}_{i^{(k)},j^{(k')}} = \alpha\sqrt{W^{(k,k')}}\exp\left[-\frac{\left\|i^{(k)} - j^{(k')}\right\|^2}{2\sigma^2}\right] \tag{5}$$

where $W^{(k,k')}_{i^{(k)},j^{(k')}}$ is the horizontal weight factor from presynaptic cells of type k' located in the j^{th} column to the postsynaptic cells of type k located in the i^{th} column, α is a normalization factor, determined numerically, which ensured that the total synaptic input integrated over all presynaptic cells of type k' to every postsynaptic cell of type k equaled $W^{(k,k')}$, σ is the Gaussian radius of the interaction, and the quantity $\left\|i^{(k)} - j^{(k')}\right\|$ denotes the horizontal distance between the pre- and post-synaptic cells, taking into account the wrap around boundary conditions employed to mitigate edge effects. An analogous weight factor describes the dependence on vertical separation. Equation 5 was augmented by a cutoff condition that prevented synaptic interactions beyond a specified distance, determined by the radius of influence of the presynaptic outputs and the postsynaptic inputs, roughly corresponding to the axonal and dendritic fields, respectively. A synaptic connection was only possible if the output radius of the presynaptic cell overlapped the input radius of the postsynaptic cell. Except for axonal connections, the input and output radii were the same for all cell types. For the large amacrine cells and the ganglion cells, the radius of influence extended out to the centers of the nearest neighboring cells of the same type, producing a coverage factor greater then one [43]. The radii of the bipolar, small, and axon-bearing amacrine cells (non-axonal connections only) extended only halfway to the nearest cell of the same type, giving a coverage factor of one [4]. The external input was multiplied by a gain factor of 3. Values for model parameters are listed in tables 1 and 2.

Figure 1 about here

Correlated Poisson Process

A second method for generating artificial spike trains employed a pseudo-Poisson processes. To produce a random spike train of duration, *T*, temporal resolution, Δt, and mean spike rate R_0, we used the formula:

$$S_i = \theta(R_0 \cdot \Delta t - r) \tag{6}$$

where $S_i = \{0,1\}$ denotes the presence or absence of a spike in the i^{th} time bin, θ is a step function, $\theta(x<0) = 0$, $\theta(x>0) = 1$, and r is a uniform random deviate between zero and one supplied by the Matlab intrinsic function RAND. In the limit that $R_0 \cdot \Delta t \ll 1$, the above procedure reduces to a Poisson process with constant rate R_0.

To produce correlated spike trains, the same time series, S_i, was used as a template to construct new spike trains according to the formula:

$$S_i^{(k)} = \theta\left(S_i \cdot C\Delta t + (1 - S_i) \cdot (1 - C\Delta t) \frac{(R_0 \Delta t)}{1 - R_0 \Delta t} - r \right) \tag{7}$$

where $S_i^{(k)}$ denotes the spike train of the k^{th} cell, C is the conditional firing rate given that a spike occurred in the corresponding time bin of the template train and r is again a uniform random deviate. The maximum value of $C \cdot \Delta t$ was $(1 - R_0 \cdot \Delta t) \cdot 0.5$.

Data Analysis

Reported correlations are expressed as a fraction of the expected synchrony due to chance, either measured during baseline activity or during the plateau portion of the response to a sustained stimulus (200-600 msec after onset). With this normalization, a correlation amplitude of one at zero delay corresponded to a doubling in the number of synchronous events over the expected rate. Correlations were plotted as a function of the delay after averaging over all events occurring during the plateau portion of the response. For each delay value, this average was compensated for edge effects arising from the finite length of the two spike trains. To increase the signal to noise, the firing rates or correlations were averaged over all cells, or distinct cell pairs, responding to the same stimulus, producing a multi-unit peri-stimulus-time-histogram (PSTH) or multi-unit cross-correlation-histogram (CCH), respectively. Auto-correlation functions were not included in the multi-unit CCH, which thus only included cross-correlations between distinct cell pairs. Error bars were estimated by assuming Poisson statistics for the count in each histogram bin. All correlations were obtained by averaging over 200 stimulus trials, using a bin width of 1 msec.

Ideal Observer

An ideal observer was used to discriminate between different intensities based on the distribution of threshold events across independent stimulus trials. For each intensity, the number of suprathreshold events was determined on successive trials and the results normalized as a probability distribution. For any given pair of intensities, the percent of correct classifications made by an ideal observer was inversely related to the degree of overlap between the two distributions. Total overlap corresponded to performance at chance (50% correct) while zero overlap implied perfect discrimination (100% correct). For the threshold detection process, all spikes occurring with a given time bin, whose width varied from 1-20 msec depending on the experiment, were summed together and the result compared to a threshold. There was no overlap between successive time bins. The duration of the recording interval, which varied from 100 to 400 msec, was adjusted so as to approximately normalize the task difficulty as the number of inputs was varied. The discrimination interval was always contained within the plateau portion of the response.

Results

Artificially Generated Spike Trains

Spike trains were generated by a computer model of the inner retina, whose parameters were adjusted to qualitatively matched recordings from cat alpha (Y) ganglion cells, which exhibit strong spatiotemporal correlations [24,29,30]. Previous theoretical studies of correlated activity have employed mathematically generated spike trains in which firing correlations were independent of, and thus conveyed no information about, the applied stimulus [25,33,34]. However, firing correlations between retinal neurons depend strongly on stimulus parameters such as size [2,16,29], contrast [29], connectedness [16,30] and velocity [29]. To generate realistic stimulus-dependent correlations, we used a detailed retinal model. Our conclusions depend only on the degree to which the firing rates and spatiotemporal correlations produced by the model are physiologically realistic, but are otherwise independent of whether we accurately identified or simulated the underlying physiological mechanisms, a separate issue addressed elsewhere [18,19].

To guard against the possibility that our mathematically generated spike trains exaggerated the information conveyed by correlations, the model was subject to the following constraints. First, when corroborating information was available in the published literature, the largest stimulus-evoked correlations between the model-generated spike trains were comparable to the correlations measured experimentally under similar conditions. Thus, under these circumstances we are confident that none of the correlations generated by the model were substantially in excess of the correlations present physiologically. Second, the level of correlations in the absence of stimulation was comparable to the spontaneous correlations measured experimentally, ensuring that a correlation code would not possess an unfair advantage by starting from an unphysiologically low baseline level. Third, because the retinal model possessed no mechanism for adaptation, we ensured that the sustained firing rates evoked by stimuli of various intensities were probably larger than those that would occur physiologically. Thus, the model was probably conservative and likely biased in favor of a rate-code.

The principal characteristics of the artificial spike trains used in this study are best illustrated by examining the multiunit-PSTHs and multiunit-CCHs computed from the simulated responses to a narrow bar stimulus of maximum intensity (fig. 2a). The multiunit-PSTH, which combined the responses of all twelve model ganglion cells directly beneath the stimulus, consisted of a sharp peak, about 30 msec wide, followed by a plateau period of sustained elevated activity (fig. 2b). Relative to baseline, the sustained increase in spike activity produced by the bar stimulus was comparable to or larger than the sustained increase in firing exhibited by cat ganglion cells in response to high contrast features [6,10]. The pronounced downward notch separating the peak and plateau regions is characteristic of ganglion cell responses to large, flashed stimuli [5].

Despite the absence of periodic structure in the multiunit-PSTH, there were nonetheless very prominent high frequency oscillations in the multiunit-CCH recorded during the plateau portion of the response (fig. 2c, solid black line). In the cat retina, high contrast stimuli produce an approximate doubling in the number of synchronous events relative to the expected level due to chance [29]. In our artificially generated spike trains, the temporal correlations produced by a maximum intensity bar stimulus exhibited a qualitatively similar increase in the number of synchronous events relative to the expected level. Thus, the artificially generated spike trains used in the present study reflect physiologically reasonable levels of correlated activity. The high frequency oscillations were not strongly phase-locked to the stimulus onset, but rather tended to become phase randomized over time, as revealed by the decline in correlation strength with increasing delay. Consistent with the absence of stimulus-locked oscillations in the model, the shift predictor was negligible (fig. 2c, dashed gray line). Correlations between cat ganglion cells exhibit qualitatively similar high frequency oscillations that decline in amplitude with increasing delay and also possess negligible shift predictors [29]. The tendency for high frequency oscillations to become phase randomized explains in part why periodic structure is not more commonly observed in PSTHs measured experimentally, as such signals tend to average out over multiple trials. Other reasons for the lack of periodic structure in experimentally recorded PSTHs include the use of large bin widths, typically 10-20 msec, and/or low pass temporal filtering, both of which are commonly employed [cf. 5,10], as well as the use of small spots or fine gratings, whereas strong oscillations are only evoked by large stimuli [2,16,29].

Figure 2 about here

Firing Correlations and Firing Rate are both Modulated by Stimulus Intensity

To investigate the relationship between correlation strength and stimulus contrast, the same narrow bar covering twelve ganglion cells was varied over a 32-fold range of intensities (fig. 3). The firing rates of the model ganglion cells activated by the stimulus, as measured by the multiunit-PSTH, increased in a graded manner as the stimulus intensity was raised (fig. 3b). Similarly, the degree of synchrony during the plateau portion of the response, assessed by the peak amplitude of the multiunit-CCH at zero delay, also increased with stimulus intensity (fig. 3b). High frequency firing correlations between cat ganglion cells depend similarly on luminance contrast [29]. As a function of stimulus intensity, synchrony could be modulated over a greater dynamic range, here measured relative to baseline activity, than could the fractional change in the multiunit firing rate (figs. 3c). Overall, firing correlations were very sensitive to stimulus intensity, and thus might convey information about contrast in additional to that represented by the mean firing rate across the ensemble of activated cells.

Figure 3 about here

Correlated Inputs Transmit More Information through Coincidence Detectors than Independent Rate-Matched Controls

A simple threshold detector with a short integration time window and a low rate of suprathreshold background events was able to extract the stimulus intensity more reliably from a hybrid rate/correlation code than from statistically independent, rate-matched Poisson distributed inputs. Spikes from a column of twelve ganglion cells activated by a narrow bar were summed into a simple threshold detector (fig. 4). The event rate of the detector was determined by the total number of times the input crossed threshold within a 200 msec epoch during the plateau portion of the response. For these experiments, the detection threshold was set to a level that required three or more spikes to arrive within the same 2 msec time bin in order to produce a detector event. To assess the extent to which the detector was able to utilize firing correlations between its inputs, the twelve stimulated ganglion cells were replaced by independent Poisson generators that, on average, produced the same number of spikes per unit time. For both correlated and Poisson input, the baseline detector event rate was very low, around 1 Hz (fig. 4a). A high-intensity stimulus produced a greater increase in the detector event rate when the inputs were correlated as compared to the case when the inputs were independent. This extra sensitivity reflected the fact that a threshold process with a short integration time is well suited for detecting rare synchronous events [17,20], a property also exhibited by cortical and sub-cortical neurons [1,41].

Figure 4 about here

When driven by correlated input from the model ganglion cells, the output of the threshold detector allowed for better discrimination between different stimulus intensities than when driven by independent Poisson generators. For each intensity, we plotted the distribution of the total number of detector events over 200 independent stimulus trials (fig. 4b). As the stimulus intensity increased, the distributions shifted to the right, reflecting the greater number of suprathreshold inputs. When the threshold detector was driven by correlated input from the retinal model, the event distributions for different intensities were more separable than when driven by the rate-matched controls. To quantify the degree to which firing correlations caused detector output to be more discriminable, we used an ideal observer based on an optimal Bayes discriminator to estimate the maximum percentage of intensity comparisons that could be correctly classified [8]. Starting from several different baselines, firing correlations allowed for a higher percentage of correct intensity classifications over a broad range of intensity increments (fig. 4c). The abscissa of each point gives the final stimulus intensity, while the x-intercept of the line passing through that point yields the corresponding baseline intensity. The ordinate of each point gives the percentage of trials on which an ideal observer could correctly classify the final stimulus intensity, relative to the baseline, using only the single trial output of the detector. For many pairwise intensity discriminations, firing correlations allowed approximately 10 additional trials out of a 100 to be correctly classified (solid black lines) compared to the Poisson control (dashed gray lines). Averaged over all intensity increments, the extra number of correct classifications was approximately 5/100, but this value is likely conservative due to saturation. Our results suggest that when retinal output is decoded by a threshold detection process with a short integration time window, firing correlations allow information to be transmitted more reliably than would be the case with independent firing rates.

Correlations Mediate Greater or Equivalent Stimulus Discrimination for a Wide Class of Threshold Detectors

The hybrid rate/correlation code continued to mediate equal or superior stimulus discrimination even when the integration time of the threshold detector was increased so as to diminish the importance of synchronous inputs. To quantify the performance of the ideal observer for a given threshold detector, we defined the quantity Δ, which gives the difference in the percentage of correctly classified trials between the hybrid rate/correlation code and the Poisson control, averaged over all intensity pairs. Our results show that a hybrid rate/correlation code yields greater or equivalent stimulus discrimination over a wide range of integration times and background activity levels (fig. 5). Plotted as a function of integration time, Δ was largest for small summation intervals well suited to resolve synchronous inputs (fig. 5a). When plotted as a function of the detection threshold, Δ generally increased for integration times less than approximately 10 msec, since larger thresholds produced lower values of background activity and thus made the detection process more sensitive to synchrony (fig. 5b).

A somewhat surprising aspect of our results was that as the integration time became large enough to effectively discard intensity information encoded by the degree of synchrony, Δ did not become strongly negative, but rather approached an asymptotic level near zero. Since spatial correlations adversely affect rate-coded information [25,33,34], it might have been anticipated that once the integration time became sufficiently long to discount any information encoded by the degree of synchrony, independent Poisson inputs would have mediated greater stimulus discrimination than correlated inputs. However, such reasoning fails to consider the effects of temporal correlations, which like spatial correlations also increased as a function of stimulus intensity, as indicated by the increased persistence of periodic structure in the multi-unit CCH as the stimulus intensity was increased (fig. 3). As a result of the stronger temporal correlations, the spike trains became more regular, and thus the total number of inputs over the course of the 200 msec trail became more reliable predictors of the stimulus intensity.

To quantify the reliability of the afferent spike trains, we computed the Fano factor of the multi-unit input as a function of stimulus intensity (fig. 5c). The Fano factor is defined as the variance in the number of spikes divided by the mean, and is equal to one for a Poisson process [39]. At all intensities, the Fano factor of the combined correlated input (solid line) was less than that of the Poisson control (dashed line), indicating that temporal correlations caused the total number of spikes to be less variable, and therefore more reliable, than for independent Poisson generators. In the absence of temporal correlations, the Fano factor would have increased markedly with stimulus intensity as a consequence of the increased spatial correlations. To illustrate this fact, a second control was employed in which stimulus-dependent spatial correlations were introduced between the separate Poisson generators. The maximum spatial correlation between the modified Poisson spike trains was set so as to approximately equal the maximum synchrony present in the hybrid rate/correlation code, while the minimum synchrony was set to zero and a linear interpolation was used for intermediate intensities. By allowing no more than one spike in each 10 msec time bin, thereby producing a relative refractory period, the Fano factor in the absence of spatial correlations could be reduced to approximately the same level exhibited by the model during background activity. In the absence of temporal correlations, strong spatial correlations between the modified Poisson generators produced a large rise in the Fano factor of the combined input as the stimulus intensity was increased (dotted line). Our results demonstrate that a refractory period, by itself, cannot account for the reliability of the total integrated input in the presence of strong spatial correlations, but that such reliability results naturally from strong temporal correlations due to high frequency oscillations. Thus, strong stimulus-dependent spatiotemporal correlations do not necessarily result in less information transfer, even when only the total spike count over a relatively long interval is considered.

Figure 5 about here

Hybrid Rate/Correlation Codes Work Best over Limited Numbers of Cells

Up to this point, we have only considered the effects of correlations between relatively small numbers of inputs. While this is consistent with the convergence ratios of sensory input onto neurons in the LGN and V1 [32,42], hybrid rate/correlation codes might cease to be advantageous as the number of inputs was increased. To explore the behavior of a hybrid rate/correlation code as a function of the number of input spike trains, we used a much larger stimulus that activated a 12×12 array of neurons. Oscillations between retinal ganglion cells increase markedly in response to the larger stimuli [2,16,29] an this was also true in our retinal model. However, because the maximum correlation strength between cat alpha cells has not been precisely determined for such stimuli, our results should be interpreted as an upper limit on the information that might be conveyed by a hybrid rate/correlation code. As a function of the number of input spike trains, Δ reached a maximum for a relatively small number of correlated inputs, at a value between 10 and 50 (fig. 6a). For integration times less than 5 msec, Δ remained greater than or equal to zero regardless of the number of inputs. For a threshold process employing longer integration times, correlations produced progressively poorer stimulus discrimination as the number of inputs increased.

As the number of inputs was raised, we increased the threshold so as to maintain the background detection rate as close to 1 Hz as possible, but without falling below 0.1 Hz. As the number of inputs increased, it was necessary to use a lower threshold for the Poisson control than for the correlated input, due to the small oscillations present in the background activity, in order to maintain the level of background suprathreshold events close to the target level of 1 Hz. We obtained similar results when the target level of background suprathreshold events was increased, but at smaller values of Δ on average. In order to maintain task difficulty as the number of inputs was increased, the duration of each trial was lowered from 400 to 100 msec. Although the actual hybrid code present physiologically may not be as robust with respect to the number of inputs as in our retinal model, the general principles illustrated by our results are nonetheless likely to be valid. A hybrid code can transmit significantly more information than a pure rate code regardless of the number of inputs as long as the target population is able to respond preferentially to synchronous events.

The dependence of the hybrid code on the number of cells feeding into the detector is paralleled by the behavior of the Fano factor of the combined input (fig. 6b). When the number of inputs was small, the Fano factor was always less than one, regardless of the stimulus intensity. As the number of inputs increased, however, the Fano factor became much greater than one at all but the smallest stimulus intensities. This result is consistent with previous studies showing that spatial correlations become progressively more destructive of rate-coded information as the number of neurons increases [25,33,34]. When the number of inputs is very large, therefore, a hybrid rate/correlation code is only likely to be effective when the integration time is small enough to resolve synchronous events.

Figure 6 about here

Discussion

Previous studies of how wide spread firing synchrony affects the representation of information in neural ensembles have focused solely on how spatial correlations impede the extraction of rate-coded signals [25,33,34]. In particular, such studies have emphasized how averaging over ensembles of similarly activated neurons only reduces the trial-to-trial variability to the extent that individual fluctuations are uncorrelated. From this point of view, spatial correlations are an inevitable feature of densely interconnected networks and as such limit the effective size of neural ensembles to several tens of strongly correlated cells. However, such analyses have generally ignored the possibility that spatiotemporal correlations themselves could encode significant information, and thus compensate for the attendant loss of rate-coded signals. Here, we have demonstrated that both spatial and temporal correlations can lead to improved information transmission under fairly general assumptions. In particular, spatial correlations, when evoked in an intensity-dependent manner, can significantly improve information transfer through threshold neurons functioning as coincidence detectors. In a complementary fashion, we found that temporal correlations, when also proportional to intensity, mediate greater stimulus discrimination by making the total number of input spikes more reliable regardless of the integration time window employed. Thus, we found that the negative effects of spatial correlations described in previous studies could be overcome in two ways: 1) by encoding similar information in the degree of synchrony and 2) by making the spike trains more regular as they become more synchronous.

General principles of population encoding and decoding, especially those involving temporal correlations between multiple neurons, are still very difficult to investigate experimentally. To reproduce the results of the present model, it would be necessary to record simultaneously from multiple input neurons and at least one common postsynaptic neuron and to manipulate the firing correlations between the input neurons either through direct multi-electrode stimulation or else via pharmacological techniques. While it is not clear that such biological experiments are currently feasible, there does exist sufficient information to construct artificial spike trains possessing physiologically realistic spatiotemporal correlations. Moreover, it is not necessary to possess a complete understanding of the physiological mechanisms that give rise to strong spatiotemporal correlations in order to analyze their information content.

While multi-electrode recording techniques allow simultaneous monitoring of tens to hundreds of neurons, such data can typically only address questions of how information is encoded across ensembles of similar cells, but not how the same information might be extracted by down stream processing elements. Here, we have focused on signal detection using a threshold process that captures many essential aspects of neural dynamics. By using a computational model, it was possible to compare the stimulus discrimination accomplished by a hybrid rate/correlation code with that mediated by a pure rate-code that produced, on average, the same number of spikes. Finally, by using a computational model, we were able to examine the effects of spatiotemporal correlations over a wide class of threshold detection processes, and in this way obtained insights into the physiological conditions necessary to utilize a hybrid rate/correlation code.

The main drawback of using a computational model is the need to ascertain whether the results are physiologically relevant. Fortunately, spatiotemporal correlations between retinal ganglion cells have been sufficiently well characterized so as to impose tight constraints on artificially generated spike trains. Where possible to verify, the model appeared to favor a rate code over a correlation code. This was the case for the narrow bar stimulus employed in our first several experiments. The maximum synchrony between the model spike trains, measured relative to chance and averaged over all cells responding to the bar, was somewhat less than the levels of synchrony between widely separated cells recorded in the cat retina in response to similar stimuli [30]. Likewise, the maximum sustained increase in firing rate in our artificial spike trains was probably exaggerated due to the lack of adaptation mechanisms in the model. Our results were also useful for illustrating general phenomena, such as how hybrid rate/correlation codes might scale as the number of cells conveying redundant information increased. The effectiveness of a hybrid rate/correlation code was found to involve a trade off between the loss of rate-coded information due to spatial correlations and the gain in information due to the information encoded by the spatial correlations themselves. At longer integration times that were insensitive to synchronous input, stimulus-dependent temporal correlations still contributed to improved performance on intensity discrimination tasks by causing spike counts to become more regular. As the number of neurons increased, the loss of rate-coded information due to spatial correlations became more severe.

Our results are likely to be helpful for interpreting spatiotemporal correlations in the retina and elsewhere in the CNS. High frequency oscillations are ubiquitous in the vertebrate retina, having been measured extracellularly in cats [22,29,30,38], rabbits [2], frogs [16] and mudpuppy [45], as well as in the ERGs of humans [7,44] and primates [13]. The conservation of retinal oscillations across such a broad range of vertebrate species suggests they may be important for visual function. Strong spatiotemporal correlations have also been recorded elsewhere in the mammalian nervous system, including visual [14,15,21,23] and sensorimotor [27,28] cortex and the hippocampus [40]. Numerous interpretations of the information processing function accomplished by spatiotemporal correlations have been suggested [9,11,12,35,37]. Here, we note that since spatiotemporal correlations are widely present throughout the brain; the nervous system might as well make use them. If correlations are indeed an unavoidable consequence of neural connectivity, our results suggests that the loss of rate-coded information could be mitigated by employing a hybrid rate/correlation code. On the other hand, correlations may also convey information that is not well represented in the simple rate-code. Our results suggest that the brain might exploit spatiotemporal correlations for higher level processing functions without sacrificing rate-coded information.

In the retina, it has been argued that information transfer is maximized when the outputs are uncorrelated [36] and recent studies indicate that the relatively strong firing correlations sometimes observed between neighboring ganglion cells convey little additional information about natural scenes [31]. Our results are not in conflict with these findings. Information theory, by itself, does not address how stimulus properties are extracted by target neurons. While formal mathematical measures predict that ganglion cells convey more information when their activity is uncorrelated, experimental and theoretical evidence suggests that synchronous inputs are particularly salient [1,17,41], and our current results demonstrate that information encoded by firing correlations between model ganglion cells can be efficiently extracted by threshold neurons under fairly general assumptions. Simultaneous recordings in cat from the retina, the LGN, and from area 18 of the visual cortex indicate that synchrony between retinal ganglion cells can be propagated to higher levels in the visual system [3]. It may therefore be necessary to consider the impact of correlated activity in order to fully account for its role in information processing. Finally, we note that using firing correlations to encode local stimulus properties does not preclude additional encoding functions that have been suggested [26].

Correspondence and requests for materials should be addressed to G.K. (e-mail: gkenyon@lanl.gov).

References:

[1] J.M. Alonso, W.M. Usrey and R.C. Reid, Precisely correlated firing in cells of the lateral geniculate nucleus, Nature 383 (1996) 815-819.

[2] M. Ariel, N.W. Daw and R.K. Rader, Rhythmicity in rabbit retinal ganglion cell responses, Vision Res. 23 (1983) 1485-1493.

[3] M. Castelo-Branco, S. Neuenschwander and W. Singer, Synchronization of visual responses between the cortex, lateral geniculate nucleus, and retina in the anesthetized cat, J Neurosci 18 (1998) 6395-6410.

[4] E. Cohen and P. Sterling, Convergence and divergence of cones onto bipolar cells in the central area of cat retina, Phil. Trans. R Soc. Lond. B 330 (1990) 323-328.

[5] J.F. Cox and M.H. Rowe, Linear and nonlinear contributions to step responses in cat retinal ganglion cells, Vision Res. 36 (1996) 2047-2060.

[6] O.D. Creutzfeldt, B. Sakmann, H. Scheich and A. Korn, Sensitivity distribution and spatial summation within receptive-field center of retinal on-center ganglion cells and transfer function of the retina, J Neurophysiol 33 (1970) 654-671.

[7] F. De Carli, L. Narici, P. Canovaro, S. Carozzo, E. Agazzi and W.G. Sannita, Stimulus- and frequency-specific oscillatory mass responses to visual stimulation in man, Clin Electroencephalogr 32 (2001) 145-151.

[8] R.O. Duda, P.E. Hart and D.G. Stork, Pattern Classification (Wiley,NY,2001).

[9] A.K. Engel, P. Fries and W. Singer, Dynamic predictions: oscillations and synchrony in top-down processing, Nat Rev Neurosci 2 (2001) 704-716.

[10] C. Enroth-Cugell and H.G. Jakiela, Suppression of cat retinal ganglion cell responses by moving patterns, J. Physiol. 302 (1980) 49-72.

[11] P. Fries, J.H. Reynolds, A.E. Rorie and R. Desimone, Modulation of oscillatory neuronal synchronization by selective visual attention, Science 291 (2001) 1560-1563.

[12] P. Fries, J.H. Schroder, P.R. Roelfsema, W. Singer and A.K. Engel, Oscillatory neuronal synchronization in primary visual cortex as a correlate of stimulus selection, J Neurosci 22 (2002) 3739-3754.

[13] L.J. Frishman, S. Saszik, R.S. Harwerth, S. Viswanathan, Y. Li, E.L. Smith, 3rd, J.G. Robson and G. Barnes, Effects of experimental glaucoma in macaques on the multifocal ERG. Multifocal ERG in laser-induced glaucoma, Doc Ophthalmol 100 (2000) 231-251.

[14] C.M. Gray, P. Konig, A.K. Engel and W. Singer, Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties, Nature 338 (1989) 334-337.

[15] C.M. Gray and W. Singer, Stimulus-specific neuronal oscillations in orientation columns of cat visual cortex, Proc Natl Acad Sci U S A 86 (1989) 1698-1702.

[16] H. Ishikane, A. Kawana and M. Tachibana, Short- and long-range synchronous activities in dimming detectors of the frog retina, Vis Neurosci 16 (1999) 1001-1014.

[17] G.T. Kenyon, E.E. Fetz and R.D. Puff. Effects of firing synchrony on signal propagation in layered networks. In: D.S. Touretzky,ed., Advances in Neural Information Processing Systems 2 (Morgan Kaufmann, San Mateo, 1990) 141-148.

[18] G.T. Kenyon, J. Jeffs, B. Moore, G.S. Stephens, K.S. Denning, B.J. Travis, J. Theiler, J.S. George and D.W. Marshak, A model of high frequency oscillations among retinal ganglion cells, Vis. Neurosci. (submitted) (submitted).

[19] G.T. Kenyon and D.W. Marshak, Gap junctions with amacrine cells provide a feedback pathway for ganglion cells within the retina, Proc R Soc Lond B Biol Sci 265 (1998) 919-925.

[20] G.T. Kenyon, R.D. Puff and E.E. Fetz, A general diffusion model for analyzing the efficacy of synaptic input to threshold neurons, Biol Cybern 67 (1992) 133-141.

[21] A.K. Kreiter and W. Singer, Stimulus-dependent synchronization of neuronal responses in the visual cortex of the awake macaque monkey, J Neurosci 16 (1996) 2381-2396.

[22] M. Laufer and M. Verzeano, Periodic activity in the visual system of the cat, Vision Res 7 (1967) 215-229.

[23] M.S. Livingstone, Oscillatory firing and interneuronal correlations in squirrel monkey striate cortex, J Neurophysiol 75 (1996) 2467-2485.

[24] D.N. Mastronarde, Correlated firing of retinal ganglion cells, Trends in Neurosciences 12 (1989) 75-80.

[25] M.E. Mazurek and M.N. Shadlen, Limits to the temporal fidelity of cortical spike rate signals, Nat Neurosci 5 (2002) 463-471.

[26] M. Meister and M.J. Berry, 2nd, The neural code of the retina, Neuron 22 (1999) 435-450.

[27] V.N. Murthy and E.E. Fetz, Oscillatory activity in sensorimotor cortex of awake monkeys: synchronization of local field potentials and relation to behavior, J Neurophysiol 76 (1996) 3949-3967.

[28] V.N. Murthy and E.E. Fetz, Synchronization of neurons during local field potential oscillations in sensorimotor cortex of awake monkeys, J Neurophysiol 76 (1996) 3968-3982.

[29] S. Neuenschwander, M. Castelo-Branco and W. Singer, Synchronous oscillations in the cat retina, Vision Res 39 (1999) 2485-2497.

[30] S. Neuenschwander and W. Singer, Long-range synchronization of oscillatory light responses in the cat retina and lateral geniculate nucleus, Nature 379 (1996) 728-732.

[31] S. Nirenberg, S.M. Carcieri, A.L. Jacobs and P.E. Latham, Retinal ganglion cells act largely as independent encoders, Nature 411 (2001) 698-701.

[32] R.C. Reid and J.M. Alonso, Specificity of monosynaptic connections from thalamus to visual cortex, Nature 378 (1995) 281-284.

[33] M.N. Shadlen and W.T. Newsome, Noise, neural codes and cortical organization, Curr Opin Neurobiol 4 (1994) 569-579.

[34] M.N. Shadlen and W.T. Newsome, The variable discharge of cortical neurons: implications for connectivity, computation, and information coding, J Neurosci 18 (1998) 3870-3896.

[35] W. Singer and C.M. Gray, Visual feature integration and the temporal correlation hypothesis, Ann. Rev. Neurosci. 18 (1995) 555-586.

[36] M.V. Srinivasan, S.B. Laughlin and A. Dubs, Predictive coding: a fresh view of inhibition in the retina, Proc R Soc Lond B Biol Sci 216 (1982) 427-459.

[37] R. Srinivasan, D.P. Russell, G.M. Edelman and G. Tononi, Increased synchronization of neuromagnetic responses during conscious perception, J Neurosci 19 (1999) 5435-5448.

[38] R.H. Steinberg, Oscillatory activity in the optic tract of cat and light adaptation, J Neurophysiol 29 (1966) 139-156.

[39] M.C. Teich, Fractal character of the auditory neural spike train, IEEE Trans Biomed Eng 36 (1989) 150-160.

[40] R.D. Traub, M.A. Whittington, S.B. Colling, G. Buzsaki and J.G. Jefferys, Analysis of gamma rhythms in the rat hippocampus in vitro and in vivo, J Physiol (Lond) 493 (1996) 471-484.

[41] W.M. Usrey, J.M. Alonso and R.C. Reid, Synaptic interactions between thalamic inputs to simple cells in cat visual cortex, J Neurosci 20 (2000) 5461-5467.

[42] W.M. Usrey, J.B. Reppas and R.C. Reid, Specificity and strength of retinogeniculate connections, J Neurophysiol 82 (1999) 3527-3540.

[43] D.I. Vaney. The mosaic of amacrine cells in the mammalian retina. In: N.N. Osborne and G.J. Chader,ed., Progress in Retinal Research (Pergamon Press, Oxford, 1990) 49-100.

[44] L. Wachtmeister, Oscillatory potentials in the retina: what do they reveal, Prog Retin Eye Res 17 (1998) 485-521.

[45] L. Wachtmeister and J.E. Dowling, The oscillatory potentials of the mudpuppy retina, IOVS 17 (1978) 1176-1188.

Figure Captions

Fig. 1. The model contained five cells types: bipolar (BP) cells, small (SA), large (LA) and poly-axonal (PA) amacrine cells, and alpha ganglion (GC) cells, arranged as a 32x32 square mosaic with wrap-around boundary conditions. Conceptually, connections could be organized into 3 categories. a) Feedforward and feedback inhibition. Excitatory synapses from BPs were balanced by a combination of reciprocal synapses and direct inhibition of the GCs, primarily mediated by the non-spiking amacrine cell types. b) Serial inhibition. The three amacrine cell types regulated each other through a negative feedback loop. c) Resonance circuit. The PAs were excited locally via electrical synapses with GCs and whose axons mediated wide field inhibition to all cell types, but most strongly onto the GCs. Note: not all connections are shown. Explanation of symbols: Excitation (triangles), inhibition (circles), gap junctions (resistors).

Fig 2. Artificial spike trains generated by the retinal model. a) A column of twelve model ganglion cells was stimulated by a narrow bar (white rectangle, intensity = ½, stimulus dimensions 2×8 GC receptive field diameters). Circles indicate GC receptive field diameter. b) Multiunit-PSTH obtained by averaging the individual PSTHs over all ganglion cells activated by the stimulus. The solid line at the bottom of the panel indicates the stimulus duration (600 msec). Vertical ticks denote the plateau portion of the response. c) Multiunit-CCH, obtained by combining individual CCHs from all distinct pairs of ganglion cells activated by the stimulus (solid black lines). Correlations expressed as a fractional change from the expected synchrony due to chance (dimensionless units). Shift predictors (dashed gray lines) obtained by recomputing the multiunit-CCH using spike trains from different stimulus trials.

Fig. 3. Firing correlations are modulated over a greater dynamic range than firing rate as a function of stimulus intensity. The stimulus was again a narrow bar covering twelve ganglion cells. a) Intensity series formed by the multiunit-PSTHs of ganglion cells activated by the stimulus. The intensity, in $\log_2$ units, is indicated at the upper right of each histogram (bin width = 10 msec). b) Intensity series formed by the multiunit-CCHs between all pairs of ganglion cells activated by the stimulus, relative to the baseline level of synchrony. Firing correlations during the plateau response are strongly modulated by stimulus intensity. c) Fractional change from baseline in synchrony (black lines, circles) as a function of stimulus intensity, compared to the fractional change in firing rate (gray lines, squares). Synchrony could be modulated over a greater dynamic range than the sustained firing rate.

Fig. 4. Firing correlations during the plateau portion of the response allow improved discrimination of stimulus intensity compared to independent Poisson input. a) Example of the threshold detection process. Stimuli consisted of a narrow bar presented at various intensities (same stimulus as in fig. 3). Ganglion cell input to the threshold detector during the plateau portion of the response is shown on the left and equivalent Poisson input on the right. The baseline activity of the detector in the absence of stimulation (top row) is very low. A stimulus with intensity = -1 in $\log_2$ units (bottom row) produced strong correlations between ganglion cells, resulting in a relatively large number of suprathreshold events. Dashed line: detector threshold. Dotted lines: average input ± std. dev. Summation window, 2 msec. b) Normalized probability distributions giving, for each stimulus intensity, the number of suprathreshold events during the 200 msec analysis interval. Left: Hybrid rate/correlation code. Right: Poisson control. The distribution of suprathreshold events produced by the hybrid code were more separable. c) Fraction of correct intensity classifications by an ideal observer. The ideal observer choose between two equally likely stimulus intensities based on the number of detector events on each trial. Input to the detector came from

either a hybrid code (solid black lines) or from Poisson controls (dashed gray lines). Each point represents the fraction of trials on which the intensity indicated by the abscissa was correctly distinguished from a lower intensity, denoted by the intersection of each line with the x-axis. Error bars computed assuming binary statistics for the overlap between each pair of distributions (omitted from Poisson controls for clarity).

Figure 5. Hybrid rate/correlation codes yield superior or equivalent stimulus discrimination for a broad class of threshold detectors. a) Δ, the difference in the percentage of successfully classified trials using a hybrid code as opposed to independent Poisson controls, plotted as a function of temporal integration window. Δ declined with increasing integration time, but did not become strongly negative. Individual points are for different thresholds. b) Δ vs. detection threshold for different integration times. Same data as in previous panel. The increase in Δ with threshold declines progressively at longer integration times. c) Fano factors (variance/mean of the total input spike count) plotted vs. stimulus intensity. Solid line+circles: Hybrid code. Dashed line+squares: Poisson control. Dotted line+triangles: Modified Poisson process in which the separate generators were correlated by an amount proportional to the stimulus intensity and no more than one spike could occur in any 10 msec time bin. The Fano factor for the hybrid code remained less than one due to temporal correlations resulting from high frequency oscillations, while spatial correlations alone produced a large increase in variability relative to independent Poisson inputs.

Figure 6. Hybrid rate/correlation codes work optimally for limited numbers of input neurons. These experiments used a 12×12 uniform spot that produced larger spatiotemporal correlations than did a narrow bar. a) Δ vs. the number of input cells plotted for several different integration times. Δ reached a peak at between 10-50 inputs and remained positive as the number of inputs increased as long as the integration time was small enough to resolve synchronous inputs, but Δ become negative for the same

number of inputs if the integration time was too long. b) Fano factor vs. number of inputs plotted for several different intensities ($\log_2$ units). At high intensities, which produce strong spatial correlations, the Fano factor increased sharply as a function of the number of inputs, thus accounting for the poorer performance of the hybrid code in this regime.

Table 1: Cellular parameters.

	τ	b	n×n	d	σ
BP	10.0	-0.0	64×64	0.25	0.25
SA	25.0	-0.5	64×64	0.25	0.25
LA	20.0	-0.25	32×32	1.0	0.5
PA	5.0	-0.025	64×64	0.25/9.0[a]	0.25/3.0[a]
GC	5.0	-0.025	32×32	1.0	0.5

Explanation of symbols: τ: time constant (msec); b: bias; n×n: array size; d: cutoff radius, σ: Gaussian radius (see eq. 5). [a]Inner radius/outer radius.

Table 2: Synaptic weights.

	BP	SA	LA	PA	GC
BP	*	-0.375[b]-	3.0b	-3.0b/-15.0c	*
SA	3.0[b]	*	-3.0[b]	0.0b/-15.0[c]	*
LA	3.0[b]	*	0.25[a]	-3.0a/-15.0[c]	*
PA	0.75[b]	-0.75b	0.25[a]	0.25a/-45.0[c]	0.25[a,d]
GC	9.0[b]	-4.5b	-4.5[b]	0.25[a]/-270.0[c]	*

Each term represents the total integrated weight from all synapses arising from the corresponding presynaptic type (columns) to each cell of the corresponding postsynaptic type (rows), (the quantity $W^{(k,k')}$ in eq. 5). Asterisk (*) indicates absence of corresponding connection. Synapse type indicated by superscript: [a]gap junction, [b]non-spiking synapse, [c]spiking synapse. [d]Maximum coupling efficiency (ratio of post- to pre-synaptic depolarization) for this gap junction synapse: DC=11.3%, Action Potential=2.7%.